On Restricted Nonnegative Matrix Factorization

Dmitry Chistikov*[1], Stefan Kiefer[2], Ines Marušić[2], Mahsa Shirmohammadi[2], and James Worrell[2]

1 **Max Planck Institute for Software Systems (MPI-SWS), Germany**
dch@mpi-sws.org
2 **University of Oxford, UK**
FirstName.LastName@cs.ox.ac.uk

arXiv:1605.07061v1 [cs.FL] 23 May 2016

Abstract

Nonnegative matrix factorization (NMF) is the problem of decomposing a given nonnegative $n \times m$ matrix M into a product of a nonnegative $n \times d$ matrix W and a nonnegative $d \times m$ matrix H. Restricted NMF requires in addition that the column spaces of M and W coincide. Finding the minimal inner dimension d is known to be NP-hard, both for NMF and restricted NMF. We show that restricted NMF is closely related to a question about the nature of minimal probabilistic automata, posed by Paz in his seminal 1971 textbook. We use this connection to answer Paz's question negatively, thus falsifying a positive answer claimed in 1974.

Furthermore, we investigate whether a rational matrix M always has a restricted NMF of minimal inner dimension whose factors W and H are also rational. We show that this holds for matrices M of rank at most 3 and we exhibit a rank-4 matrix for which W and H require irrational entries.

1998 ACM Subject Classification F.1.1 Models of Computation, F.2.1 Numerical Algorithms and Problems

Keywords and phrases nonnegative matrix factorization, nonnegative rank, probabilistic automata, labelled Markov chains, minimization

Digital Object Identifier 10.4230/LIPIcs.ICALP.2016.XXX

1 Introduction

Nonnegative matrix factorization (NMF) is the task of factoring a matrix of nonnegative real numbers M (henceforth a *nonnegative* matrix) as a product $M = W \cdot H$ such that matrices W and H are also nonnegative. The smallest inner dimension of any such factorization is called the *nonnegative rank* of M, written $\mathrm{rank}_+(M)$.

In machine learning, NMF was popularized by the seminal work of Lee and Seung [14] as a tool for finding features in facial-image databases. Since then, NMF has found a broad range of applications—including document clustering, topic modelling, computer vision, recommender systems, bioinformatics, and acoustic signal processing [5, 4, 7, 19, 21, 22]. In applications, matrix M can typically be seen as a matrix of data points: each column of M corresponds to a data point and each row to a feature. Then, computing a nonnegative factorization $M = W \cdot H$ corresponds to expressing the data points (columns of M) as convex combinations of latent factors (columns of W), i.e., as linear combinations of latent factors with nonnegative coefficients (columns of H).

* Present address: Department of Computer Science, University of Oxford, UK.

43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016).
Editors: Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi; Article No. XXX; pp. XXX:1–XXX:24
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

From a computational perspective, perhaps the most basic problem concerning NMF is whether a given nonnegative matrix of rational numbers M admits an NMF with inner dimension at most a given number k. Formally, the *NMF problem* asks whether $\text{rank}_+(M) \leq k$. In practical applications, various heuristics and local-search algorithms are used to compute an approximate nonnegative factorization, but little is known in terms of their theoretical guarantees. The NMF problem under the *separability* assumption of Donoho and Stodden [9] is tractable: an NMF $M = W \cdot H$ is called *separable* if every column of W is also a column of M. In 2012, Arora et al. [2] showed that it is decidable in polynomial time whether a given matrix admits a separable NMF with a given inner dimension. Further progress was made recently, with several efficient algorithms for computing near-separable NMFs [13, 12].

Vavasis [20] showed that the problem of deciding whether the rank of a nonnegative matrix is equal to its nonnegative rank is NP-hard. This result implies that generalizations of this problem, such as the aforementioned NMF problem, the problem of computing the factors W, H (in both exact and approximate versions), and nonnegative rank determination, are also NP-hard. It is not known whether any of these problems are in NP.

Vavasis [20] notes that the difficulty in proving membership in NP lies in the fact that a certificate for a positive answer to the NMF problem seems to require the sought factors: a pair of nonnegative matrices W, H such that $M = W \cdot H$. Related to this, Cohen and Rothblum [8] posed the question of whether, given a nonnegative matrix of rational numbers M, there always exists an NMF $M = W \cdot H$ of inner dimension equal to $\text{rank}_+(M)$ such that both W and H are also matrices of rational numbers. A natural route to proving membership of the NMF problem in NP would be to give a positive answer to the question of Cohen and Rothblum (as well as a polynomial bound on the bit-length of the factors W and H). However, the question remains open. Currently the best complexity bound for the NMF problem is membership in PSPACE, which is obtained by translation into the existential theory of real-closed fields [2]. Such a translation shows that one can always choose the entries of W and H to be algebraic numbers.

In this work, we focus on the so-called *restricted* NMF (RNMF) problem, introduced by Gillis and Glineur [11]. The RNMF problem is defined as the NMF problem, except that the column spaces of M and W are required to coincide. (Note that for any NMF, the column space of M is a subspace of the column space of W.) This problem has a natural geometric interpretation as the *nested polytope problem (NPP)*: the problem of finding a minimum-vertex polytope nested between two given convex polytopes. In more detail, for a rank-r matrix M, finding an RNMF with inner dimension d is known to correspond exactly to finding a nested polytope with d vertices in an $(r-1)$-dimensional NPP.

Our contributions are as follows.

1. We establish a tight connection between NMF and the coverability relation in labelled Markov chains (LMCs). The latter notion was introduced by Paz [15]. Loosely speaking, an LMC $\mathcal{M}'$ *covers* an LMC $\mathcal{M}$ if for any initial distribution over the states of $\mathcal{M}$ there is an initial distribution over the states of $\mathcal{M}'$ such that $\mathcal{M}$ and $\mathcal{M}'$ are equivalent. In 1971, Paz [15] asked a question about the nature of minimal covering LMCs. The question was supposedly answered positively in 1974 [3]. However, we show that the correct answer is negative, thus falsifying the claim in [3]. Instrumental to our counterexample is the observation that restricted nonnegative rank and nonnegative rank can be different. (Indeed, the wrong claims in [3] seem to implicitly rely on the opposite assumption, although the notions of NMF and RNMF had not yet been developed.)
2. We show that the RNMF problem for matrices M of rank 3 or less can be solved in polynomial time. In fact, we show that there is always a *rational* NMF of M with inner

dimension $\mathrm{rank}_+(M)$, and that it can be computed in polynomial time in the Turing model of computation. This improves a result in [11] where the RNMF problem is shown to be solvable in polynomial time assuming a RAM model with unit-cost arithmetic. Both our algorithm and the one in [11] exploit the connection to the 2-dimensional NPP, allowing us to take advantage of a geometric algorithm by Aggarwal et al. [1]. We need to adapt the algorithm in [1] to ensure that the occurring numbers are rational and can be computed in polynomial time in the Turing model of computation.

3. We exhibit a rank-4 matrix that has an RNMF with inner dimension 5 but *no rational* RNMF with inner dimension 5. We construct this matrix via a particular instance of the 3-dimensional NPP, again taking advantage of the geometric interpretation of RNMF. Our result answers the RNMF variant of Cohen and Rothblum's question in [8] negatively. The original (NMF) variant remains open.

Detailed proofs of all results can be found in the appendix.

2 Nonnegative Matrix Factorization

Let $\mathbb{N}$ and $\mathbb{N}_0$ denote the set of all positive and nonnegative integers, respectively. For every $n \in \mathbb{N}$, we write $[n]$ for the set $\{1, 2, \ldots, n\}$ and write I_n for the identity matrix of order n. For any ordered field $\mathbb{F}$, we denote by $\mathbb{F}_+$ the set of all its nonnegative elements. For any vector v, we write v_i for its i^{th} entry. A vector v is called *stochastic* if its entries are nonnegative real numbers that sum up to one. For every $i \in [n]$, we write e_i for the i^{th} coordinate vector in $\mathbb{R}^n$. We write $\mathbf{1}^{(n)}$ for the n-dimensional column vector with all ones. We omit the superscript if it is understood from the context.

For any matrix M, we write M_i for its i^{th} row, M^j for its j^{th} column, and $M_{i,j}$ for its $(i,j)^{\text{th}}$ entry. The *column space* (resp., *row space*) of M, written $\mathrm{Col}(M)$ (resp., $\mathrm{Row}(M)$), is the vector space spanned by the columns (resp., rows) of M. A matrix is called *nonnegative* (resp., *zero* or *rational*) if so are all its entries. A nonnegative matrix is *column-stochastic* (resp., *row-stochastic*) if the element sum of each of its columns (resp., rows) is one.

2.1 Nonnegative Rank

Let $\mathbb{F}$ be an ordered field, such as the reals $\mathbb{R}$ or the rationals $\mathbb{Q}$. Given a nonnegative matrix $M \in \mathbb{F}_+^{n\times m}$, a *nonnegative matrix factorization (NMF) over* $\mathbb{F}$ of M is any representation of the form $M = W \cdot H$ where $W \in \mathbb{F}_+^{n\times d}$ and $H \in \mathbb{F}_+^{d\times m}$ are nonnegative matrices. Note that $\mathrm{Col}(M) \subseteq \mathrm{Col}(W)$. We refer to d as the *inner dimension* of the NMF, and hence refer to NMF $M = W \cdot H$ as being *d-dimensional*. The *nonnegative rank over* $\mathbb{F}$ of M is the smallest number $d \in \mathbb{N}_0$ such that there exists a d-dimensional NMF over $\mathbb{F}$ of M. An equivalent characterization [8] of the nonnegative rank over $\mathbb{F}$ of M is as the smallest number of rank-1 matrices in $\mathbb{F}_+^{n\times m}$ such that M is equal to their sum. The nonnegative rank over $\mathbb{R}$ will henceforth simply be called nonnegative rank, and will be denoted by $\mathrm{rank}_+(M)$. For any nonnegative matrix $M \in \mathbb{R}_+^{n\times m}$, it is easy to see that $\mathrm{rank}(M) \le \mathrm{rank}_+(M) \le \min\{n, m\}$.

Given a nonzero matrix $M \in \mathbb{F}_+^{n\times m}$, by removing the zero columns of M and dividing each remaining column by the sum of its elements, we obtain a column-stochastic matrix M' with equal nonnegative rank. Similarly, if $M = W \cdot H$ then after removing zero columns in W and multiplying with a suitable diagonal matrix D, we get $M = W \cdot H = WD \cdot D^{-1}H$ where WD is column-stochastic. If M is column-stochastic then $\mathbf{1}^\top = \mathbf{1}^\top M = \mathbf{1}^\top WD \cdot D^{-1}H = \mathbf{1}^\top D^{-1}H$, hence $D^{-1}H$ is column-stochastic as well. Thus, without loss of generality one can consider NMFs of column-stochastic matrices into column-stochastic matrices [8, Theorem 3.2].

NMF problem: Given a matrix $M \in \mathbb{Q}_+^{n\times m}$ and $k \in \mathbb{N}$, is $\operatorname{rank}_+(M) \leq k$?

The NMF problem is NP-hard, even for $k = \operatorname{rank}(M)$ (see [20]). On the other hand, it is reducible to the existential theory of the reals, hence by [6, 16] it is in PSPACE.

For a matrix $M \in \mathbb{Q}_+^{n\times m}$, its nonnegative rank over $\mathbb{Q}$ is clearly at least $\operatorname{rank}_+(M)$. While those ranks are equal if $\operatorname{rank}(M) \leq 2$, a longstanding open question by Cohen and Rothblum asks whether they are always equal [8]. In other words, it is conceivable that there exists a rational matrix $M \in \mathbb{Q}_+^{n\times m}$ with $\operatorname{rank}_+(M) = d$ that has no *rational* NMF with inner dimension d. Recently, Shitov [17] exhibited a nonnegative matrix (with irrational entries) whose nonnegative rank over a subfield of $\mathbb{R}$ is different from its nonnegative rank over $\mathbb{R}$.

2.2 Restricted Nonnegative Rank

For all matrices $M \in \mathbb{F}_+^{n\times m}$, an NMF $M = W \cdot H$ is called *restricted NMF (RNMF)* [11] if $\operatorname{rank}(M) = \operatorname{rank}(W)$. As we know $\operatorname{Col}(M) \subseteq \operatorname{Col}(W)$ holds for all NMF instances, the condition $\operatorname{rank}(M) = \operatorname{rank}(W)$ is then equivalent to $\operatorname{Col}(M) = \operatorname{Col}(W)$. The *restricted nonnegative rank over* $\mathbb{F}$ of M is the smallest number $d \in \mathbb{N}_0$ such that there exists a d-dimensional restricted nonnegative factorization over $\mathbb{F}$ of M. Unless indicated otherwise, henceforth we will assume $\mathbb{F} = \mathbb{R}$ when speaking of the restricted nonnegative rank of M, and denote it by $\operatorname{rrank}_+(M)$.

RNMF problem: Given a matrix $M \in \mathbb{Q}_+^{n\times m}$ and $k \in \mathbb{N}$, is $\operatorname{rrank}_+(M) \leq k$?

We have the following basic properties.

▶ **Lemma 1** ([11]). *Let* $M \in \mathbb{R}_+^{n\times m}$. *Then* $\operatorname{rank}(M) \leq \operatorname{rank}_+(M) \leq \operatorname{rrank}_+(M) \leq m$. *Moreover, if* $\operatorname{rank}(M) = \operatorname{rank}_+(M)$ *then* $\operatorname{rank}(M) = \operatorname{rrank}_+(M)$.

Thus, with the above-mentioned NP-hardness result, it follows that the RNMF problem is also NP-hard and in PSPACE.

For a matrix $M \in \mathbb{Q}_+^{n\times m}$, its restricted nonnegative rank over $\mathbb{Q}$ is clearly at least $\operatorname{rrank}_+(M)$. As with nonnegative rank, in general it is not known whether the restricted nonnegative ranks of M over $\mathbb{R}$ and over $\mathbb{Q}$ are equal. By [8, Theorem 4.1] and Lemma 1, this is true when $\operatorname{rank}(M) \leq 2$.

RNMF has the following geometric interpretation. For a dimension $\ell \in \mathbb{N}$, the *convex combination* of a set $\{v_1, \ldots, v_m\} \subset \mathbb{R}^\ell$ is a point $\lambda_1 v_1 + \cdots + \lambda_m v_m$ where $(\lambda_1, \ldots, \lambda_m)$ is a stochastic vector. The *convex hull* of $\{v_1, \ldots, v_m\}$, written as $\operatorname{conv}\{v_1, \ldots, v_m\}$, is the set of all convex combinations of $\{v_1, \ldots, v_m\}$. We call $\operatorname{conv}\{v_1, \ldots, v_m\}$ a *polytope spanned by* $v_1, \ldots, v_m$. A *polyhedron* is a set $\{\, x \in \mathbb{R}^\ell \mid Ax + b \geq 0 \,\}$ with $A \in \mathbb{R}^{n\times\ell}$ and $b \in \mathbb{R}^n$. A set is a polytope if and only if it is a bounded polyhedron. A polytope is *full-dimensional* (i.e., has volume) if the matrix $(A \;\; b) \in \mathbb{R}^{n\times(\ell+1)}$ has rank $\ell + 1$.

Nested polytope problem (NPP): Given $r, n \in \mathbb{N}$, let $A \in \mathbb{Q}^{n\times(r-1)}$ and $b \in \mathbb{Q}^n$ be such that $P = \{\, x \in \mathbb{R}^{r-1} \mid Ax + b \geq 0 \,\}$ is a full-dimensional polytope. Let $S \subseteq P$ be a full-dimensional polytope described by spanning points. The *nested polytope problem (NPP)* asks, given A, b, S and a number $k \in \mathbb{N}$, whether there exist k points that span a polytope Q with $S \subseteq Q \subseteq P$. Such a polytope Q is called *nested* between P and S.

The following proposition appears as Theorem 1 in [11].

▶ **Proposition 2.** *The RNMF problem and the NPP are interreducible in polynomial time.*

More specifically, the reductions are as follows.

1. Given a nonnegative matrix $M \in \mathbb{Q}_+^{n\times m}$ of rank r, one can compute in polynomial time $A \in \mathbb{Q}^{n\times(r-1)}$ and $b \in \mathbb{Q}^n$ such that $P = \{\, x \in \mathbb{R}^{r-1} \mid Ax + b \geq 0 \,\}$ is a full-dimensional polytope, and m rational points that span a full-dimensional polytope $S \subseteq P$ such that
   (a) any d-dimensional RNMF (rational or irrational) of M determines d points that span a polytope Q with $S \subseteq Q \subseteq P$, and
   (b) any d points (rational or irrational) that span a polytope Q with $S \subseteq Q \subseteq P$ determine a d-dimensional RNMF of M.
2. Let $A \in \mathbb{Q}^{n\times(r-1)}$ and $b \in \mathbb{Q}^n$ such that $P = \{\, x \in \mathbb{R}^{r-1} \mid Ax+b \geq 0 \,\}$ is a full-dimensional polytope. Let $S \subseteq P$ be a full-dimensional polytope spanned by $s_1, \ldots, s_m \in \mathbb{Q}^{r-1}$. Then matrix $M \in \mathbb{Q}^{n\times m}$ with $M^i = As_i + b$ for $i \in [m]$ satisfies (a) and (b).

Importantly, the correspondences (a) and (b) preserve rationality. In Appendix A we detail the reduction from point 2 above, thereby filling in a small gap in the proof of [11].

▶ **Example 3** ([11, Example 1])**.** Using the geometric interpretation of restricted nonnegative rank it follows easily that, in general, we may have $\operatorname{rank}(M) < \operatorname{rank}_+(M) < \operatorname{rrank}_+(M)$. Let *3D-cube* NPP be the NPP instance where the inner and outer polytope are the standard 3D cube, i.e., $P = S = \{\, x \in \mathbb{R}^3 \mid x_i \in [0,1],\ 1 \leq i \leq 3 \,\}$. The only nested polytope is $Q = P$. The corresponding restricted NMF problem consists of the following matrix $M \in \mathbb{R}_+^{6\times 8}$:

$$M = \begin{pmatrix} 0&0&0&0&1&1&1&1 \\ 1&1&1&1&0&0&0&0 \\ 0&0&1&1&0&0&1&1 \\ 1&1&0&0&1&1&0&0 \\ 0&1&0&1&0&1&0&1 \\ 1&0&1&0&1&0&1&0 \end{pmatrix}.$$

We have $\operatorname{rrank}_+(M) = 8$ and $\operatorname{rank}(M) = 4$. Since $\operatorname{rank}_+(M)$ is bounded above by the number of rows in M, we have $\operatorname{rank}_+(M) \leq 6$. It is shown in [11] that $\operatorname{rank}_+(M) = 6$.

3 Coverability of Labelled Markov Chains

In this section, we establish a connection between RNMF and the coverability relation for labelled Markov chains. We thereby answer an open question posed in 1971 by Paz [15] about the nature of minimal covering labelled Markov chains.

A *labelled Markov chain* (*LMC*) is a tuple $\mathcal{M} = (n, \Sigma, \mu)$ where $n \in \mathbb{N}$ is the number of states, Σ is a finite alphabet of labels, and function $\mu : \Sigma \to [0,1]^{n\times n}$ specifies the transition matrices and is such that $\sum_{\sigma\in\Sigma} \mu(\sigma)$ is a row-stochastic matrix. The intuitive behaviour of the LMC $\mathcal{M}$ is as follows: When $\mathcal{M}$ is in state $i \in [n]$, it emits label σ and moves to state j, with probability $\mu(\sigma)_{i,j}$.

We extend the function μ to words by defining $\mu(\varepsilon) := I_n$ and $\mu(\sigma_1 \ldots \sigma_k) := \mu(\sigma_1) \cdots \mu(\sigma_k)$ for all $k \in \mathbb{N}$, and all $\sigma_1, \ldots, \sigma_k \in \Sigma$. Observe that $\mu(xy) = \mu(x) \cdot \mu(y)$ for all words $x, y \in \Sigma^*$. We view $\mu(w)$ for a word $w \in \Sigma^*$ as follows: if $\mathcal{M}$ is in state $i \in [n]$, it emits w and moves to state j in $|w|$ steps, with probability $\mu(w)_{i,j}$.

For $i \in [n]$ and $w \in \Sigma^*$, we write $pr_i^{\mathcal{M}}(w) := e_i^\top \cdot \mu(w) \cdot \mathbf{1}^{(n)}$ for the probability that, starting in state i, $\mathcal{M}$ emits word w. For example, in Figure 1 we have $pr_0^{\mathcal{M}}(a_1 b_1) = \frac{1}{12}$. More generally, for a given *initial distribution* π on the set of states $[n]$ (viewed as a stochastic row vector), we write $pr_\pi^{\mathcal{M}}(w) := \pi \cdot \mu(w) \cdot \mathbf{1}^{(n)}$ for the probability that $\mathcal{M}$ emits word w starting from state distribution π. We omit the superscript $\mathcal{M}$ from $pr_\pi^{\mathcal{M}}$ when it is clear from the context.

We say that an LMC $\mathcal{M}$ is *covered by* an LMC $\mathcal{M}'$, written as $\mathcal{M}' \geq \mathcal{M}$, if for every initial distribution π for $\mathcal{M}$ there exists a distribution π' for $\mathcal{M}'$ such that $pr_\pi^{\mathcal{M}}(w) = pr_{\pi'}^{\mathcal{M}'}(w)$ for all words $w \in \Sigma^*$.

The *backward matrix* of $\mathcal{M}$ is a matrix $\operatorname{Back}\mathcal{M} \in \mathbb{R}_+^{[n]\times\Sigma^*}$ where $(\operatorname{Back}\mathcal{M})_{i,w} = pr_i^{\mathcal{M}}(w)$ for every $i \in [n]$ and $w \in \Sigma^*$. The *rank* of $\mathcal{M}$ is defined by $\operatorname{rank}(\mathcal{M}) = \operatorname{rank}(\operatorname{Back}\mathcal{M})$. (Matrix $\operatorname{Back}\mathcal{M}$ is infinite, but since it has n rows, its rank is at most n.) It follows easily from the definition (see also [15, Theorem 3.1]) that $\mathcal{M}' \geq \mathcal{M}$ if and only if there exists a row-stochastic matrix A such that $A \cdot \operatorname{Back}\mathcal{M}' = \operatorname{Back}\mathcal{M}$.

LMCs can be seen as a special case of stochastic sequential machines, a class of probabilistic automata introduced and studied by Paz [15]. More specifically, they are stochastic sequential machines with a singleton input alphabet and Σ as output alphabet. In his seminal 1971 textbook on probabilistic automata [15], Paz asks the following question:

▶ **Question 4** (Paz [15], p. 38)**.** *If an n-state LMC $\mathcal{M}$ is covered by an n'-state LMC $\mathcal{M}'$ where $n' < n$, is $\mathcal{M}$ necessarily covered by some n^*-state LMC $\mathcal{M}^*$, where $n^* < n$, such that $\mathcal{M}^*$ and $\mathcal{M}$ have the same rank?*

In 1974, a positive answer to this question was claimed [3, Theorem 13]. In fact, the author of [3] makes a stronger claim, namely that the answer to Question 4 is yes, even if the inequality $n^* < n$ in Question 4 is replaced by $n^* \leq n'$. To the contrary, we show:

▶ **Theorem 5.** *The answer to Question 4 is negative.*

Theorem 5 falsifies the claim in [3]. In Appendix B.1 we discuss in detail the mistake in [3]. To prove Theorem 5 we establish a tight connection between NMF and LMC coverability:

▶ **Proposition 6.** *Given a nonnegative matrix $M \in \mathbb{Q}_+^{n\times m}$ of rank r, one can compute in polynomial time an LMC $\mathcal{M} = (m+2, \Sigma, \mu)$ of rank $r+2$ such that for all $d \in \mathbb{N}$:*

(a) any d-dimensional NMF $M = W \cdot H$ determines an LMC $\mathcal{M}' = (d+2, \Sigma, \mu')$ with $\mathcal{M}' \geq \mathcal{M}$ and $\operatorname{rank}(\mathcal{M}') = \operatorname{rank}(W) + 2$, and

(b) any LMC $\mathcal{M}' = (d+2, \Sigma, \mu')$ with $\mathcal{M}' \geq \mathcal{M}$ determines a d-dimensional NMF $M = W \cdot H$ with $\operatorname{rank}(\mathcal{M}') = \operatorname{rank}(W) + 2$.

In particular, for all $d \in \mathbb{N}$ the inequality $\operatorname{rrank}_+(M) \leq d$ holds if and only if $\mathcal{M}$ is covered by some $(d+2)$-state LMC $\mathcal{M}'$ such that $\mathcal{M}'$ and $\mathcal{M}$ have the same rank.

Assuming Proposition 6 we can prove Theorem 5:

Proof of Theorem 5. Let $M \in \{0,1\}^{6\times 8}$ be the matrix from Example 3. Let $\mathcal{M} = (10, \Sigma, \mu)$ be the associated LMC from Proposition 6. Since $M = I_6 \cdot M$ is an NMF with inner dimension 6, by Proposition 6 (a) there is an LMC $\mathcal{M}' = (8, \Sigma, \mu')$ with $\mathcal{M}' \geq \mathcal{M}$. Towards a contradiction, suppose the answer to Question 4 were yes. Then $\mathcal{M}$ is also covered by some n^*-state LMC $\mathcal{M}^*$, where $n^* \leq 9$, such that $\mathcal{M}^*$ and $\mathcal{M}$ have the same rank. The last sentence of Proposition 6 then implies that $\operatorname{rrank}_+(M) \leq 7$. But this contradicts the equality $\operatorname{rrank}_+(M) = 8$ from Example 3. Hence, the answer to Question 4 is no. ◀

To prove Proposition 6 we adapt a reduction from NMF to the trace-refinement problem in Markov decision processes [10].

Proof sketch of Proposition 6. Let $M \in \mathbb{Q}_+^{n\times m}$ be a nonnegative matrix of rank r. As argued in Section 2.1, without loss of generality we may assume that M is column-stochastic and consider factorizations of M into column-stochastic matrices only.

We define an LMC $\mathcal{M} = (m+2, \Sigma, \mu)$ with $m+2$ states $\{0, 1, \ldots, m, m+1\}$. The alphabet is $\Sigma = \{a_1, \ldots, a_m\} \cup \{b_1, \ldots, b_n\} \cup \{\checkmark\}$ and the function μ, for all $i \in [m]$ and all $j \in [n]$, is defined by:

$$\mu(a_i)_{0,i} = \tfrac{1}{m}, \qquad \mu(b_j)_{i,m+1} = (M^\top)_{i,j} = M_{j,i}, \qquad \mu(\checkmark)_{m+1,m+1} = 1,$$

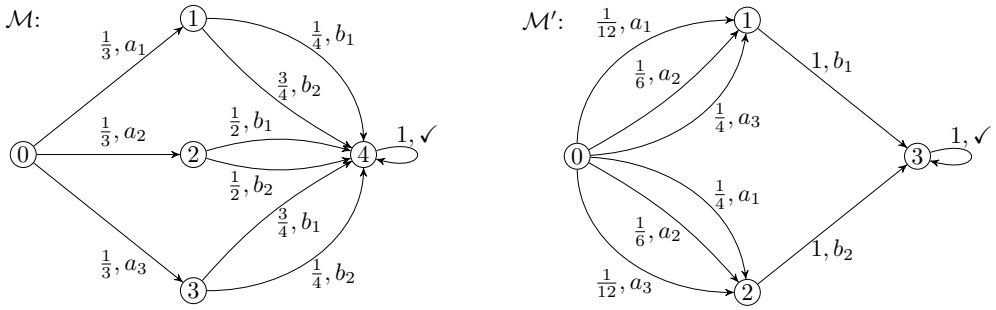

Figure 1 LMC $\mathcal{M}$ is constructed from matrix $M = \left(\begin{smallmatrix} 1/4 & 1/2 & 3/4 \\ 3/4 & 1/2 & 1/4 \end{smallmatrix}\right)$ whereas LMC $\mathcal{M}'$ is obtained by NMF $M = I_2 \cdot M$.

and all other entries of $\mu(a_i)$, $\mu(b_j)$, and $\mu(\checkmark)$ are 0. See Figure 1 for an example. We have:

$$\operatorname{Back}\mathcal{M} = \begin{array}{c} \begin{matrix} \varepsilon & b_1 & \cdots & b_n & \checkmark & a_i & a_ib_j & b_1\checkmark & \cdots & b_n\checkmark & \checkmark^2 & \cdots \end{matrix} \\ \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 & \frac{1}{m} & \frac{1}{m}M_{j,i} & 0 & \cdots & 0 & 0 & \cdots \\ 1 & M_{1,1} & \cdots & M_{n,1} & 0 & 0 & 0 & M_{1,1} & \cdots & M_{n,1} & 0 & \cdots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \cdots \\ 1 & M_{1,m} & \cdots & M_{n,m} & 0 & 0 & 0 & M_{1,m} & \cdots & M_{n,m} & 0 & \cdots \\ 1 & 0 & \cdots & 0 & 1 & 0 & 0 & 0 & \cdots & 0 & 1 & \cdots \end{pmatrix} \end{array}.$$

Thus $\operatorname{rank}(\mathcal{M}) = \operatorname{rank}(\operatorname{Back}\mathcal{M}) \geq \operatorname{rank}(M) + 2$. The first $n+2$ columns (indexed by $\varepsilon, b_1, \ldots, b_n, \checkmark$) in $\operatorname{Back}\mathcal{M}$ span $\operatorname{Col}(\operatorname{Back}\mathcal{M})$. Therefore, $\operatorname{rank}(\mathcal{M}) = \operatorname{rank}(M) + 2 = r + 2$.

For $d \in \mathbb{N}$, let $M = W \cdot H$ for some column-stochastic matrices $W \in \mathbb{R}_+^{n\times d}$ and $H \in \mathbb{R}_+^{d\times m}$. Define an LMC $\mathcal{M}' = (d+2, \Sigma, \mu')$ where the states are $\{0, 1, \ldots, d, d+1\}$. The function μ', for all $i \in [m]$, $j \in [n]$, and $l \in [d]$, is defined by:

$$\mu'(a_i)_{0,l} = \tfrac{1}{m}H_{l,i}, \qquad \mu'(b_j)_{l,d+1} = W_{j,l}, \qquad \mu'(\checkmark)_{d+1,d+1} = 1,$$

and all other entries of $\mu'(a_i)$, $\mu'(b_j)$, and $\mu'(\checkmark)$ are 0. From the NMF $M = W \cdot H$ it follows that we can factor $\operatorname{Back}\mathcal{M}$ as follows:

$$\begin{pmatrix} 1 & 0 & \cdots & 0 & 0 \\ 0 & H_{1,1} & \cdots & H_{d,1} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & H_{1,m} & \cdots & H_{d,m} & 0 \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix} \cdot \begin{array}{c} \begin{matrix} \varepsilon & b_1 & \cdots & b_n & \checkmark & a_i & a_ib_j & b_j\checkmark & \checkmark^2 & \cdots \end{matrix} \\ \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 & \frac{1}{m} & \frac{1}{m}M_{j,i} & 0 & 0 & \cdots \\ 1 & W_{1,1} & \cdots & W_{n,1} & 0 & 0 & 0 & W_{j,1} & 0 & \cdots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \cdots \\ 1 & W_{1,d} & \cdots & W_{n,d} & 0 & 0 & 0 & W_{j,d} & 0 & \cdots \\ 1 & 0 & \cdots & 0 & 1 & 0 & 0 & 0 & 1 & \cdots \end{pmatrix} \end{array}$$

where the left factor is row-stochastic (as H is column-stochastic), and the right factor equals $\operatorname{Back}\mathcal{M}'$. It follows that $\mathcal{M}' \geq \mathcal{M}$. ◀

4 Restricted NMF of Rank-3 Matrices

In this section we consider rational matrices of rank at most 3. We show that for such matrices the restricted nonnegative ranks over $\mathbb{R}$ and $\mathbb{Q}$ are equal and we give a polynomial-time algorithm that computes a minimal-dimension RNMF over $\mathbb{Q}$.

▶ **Theorem 7.** *Given a matrix $M \in \mathbb{Q}_+^{n\times m}$ where $\operatorname{rank}(M) \leq 3$, there is a rational RNMF of M with inner dimension $\operatorname{rrank}_+(M)$ and it can be computed in polynomial time in the Turing model of computation.*

Using reduction 1 of Proposition 2, we can reduce in polynomial time the RNMF problem for rank-3 matrices to the 2-dimensional NPP, i.e., the nested polygon problem in the plane. As noted in Section 2.2, the correspondence between restricted nonnegative factorizations and nested polygons preserves rationality. Thus to prove Theorem 7 it suffices to prove:

▶ **Theorem 8.** *Given polygons $S \subseteq P \subseteq \mathbb{R}^2$ with rational vertices, there exists a minimum-vertex polygon Q nested between P and S that also has rational vertices. Moreover there is an algorithm that, given P and S, computes such a polygon in polynomial time in the Turing machine model.*

In fact, Aggarwal et al. [1] give an algorithm for the 2-dimensional NPP and prove that it runs in polynomial time in the RAM model with unit-cost arithmetic. However, they freely use trigonometric functions and do not address the rationality of the output of the algorithm nor its complexity in the Turing model. To prove Theorem 8 we show that, by adopting a suitable representation of the vertices of a nested polygon, the algorithm in [1] can be adapted so that it runs in polynomial time in the Turing model. We furthermore use this representation to prove that the minimum-vertex nested polygon identified by the resulting algorithm has rational vertices.

The remainder of the section is devoted to the proof of Theorem 8. We first recall some terminology from [1] and describe their algorithm.

A *supporting line segment* is a directed line segment, with its initial and final points on the boundary of the outer polygon P, that touches the inner polygon S on its left. A nested polygon with vertices on the boundary of P is said to be *supporting* if all but at most one of its edges are supporting line segments. A polygon nested between P and S is called *minimal* if it has the minimum number of vertices among all polygons nested between P and S. It is shown in [1, Lemma 4] that there is always a supporting polygon that is also minimal, and the algorithm given therein outputs such a polygon.

Let k be the number of vertices of a minimal nested polygon. Given a vertex v on the boundary of P, there is a uniquely defined supporting polygon Q_v with at most $k+1$ vertices. To determine Q_v one computes the supporting line segments $v_1v_2, \ldots, v_kv_{k+1}$, where $v_1 = v$; see Figure 2. Then Q_v is either the polygon with vertices $v_1, \ldots, v_k$ or the polygon with vertices $v_1, \ldots, v_{k+1}$. In the first case, Q_v is minimal. The idea behind the algorithm of [1] is to search along the boundary of P for an initial vertex v such that Q_v is minimal.

As a central ingredient for our proof of Theorem 8, we choose a convenient representation of the vertices of supporting polygons. To this end, we assume that the edges of P are oriented counter-clockwise, and we represent a vertex v on an edge pq of P by the unique $\lambda \in [0, 1]$ such that $v = (1 - \lambda)p + \lambda q$. We call this the *convex representation* of v.

Similar to [1], we associate with each supporting line segment uv a *ray function* r, such that if λ is the convex representation of u then $r(\lambda)$ is the convex representation of v. The same ray function applies for supporting line segments $u'v'$ with u' in a small enough interval containing u.

In the following, when we say polynomial time, we mean polynomial time in the Turing model. To obtain a polynomial time bound, the key lemma is as follows:

▶ **Lemma 9.** *Consider bounded polygons $S \subseteq P \subseteq \mathbb{R}^2$ whose vertices are rational and of bit-length L. Then the ray function associated with a supporting line segment uv has the form $r(\lambda) = \frac{a\lambda+b}{c\lambda+d}$, where coefficients a, b, c, d are rational numbers with bit-length $O(L)$ that can be computed in polynomial time.*

Suppose that $v_1v_2, \ldots, v_kv_{k+1}$ is a sequence of k supporting line segments, with corresponding ray functions $r_1, \ldots, r_k$. Then $v_1, \ldots, v_k$ are the vertices of a minimal supporting

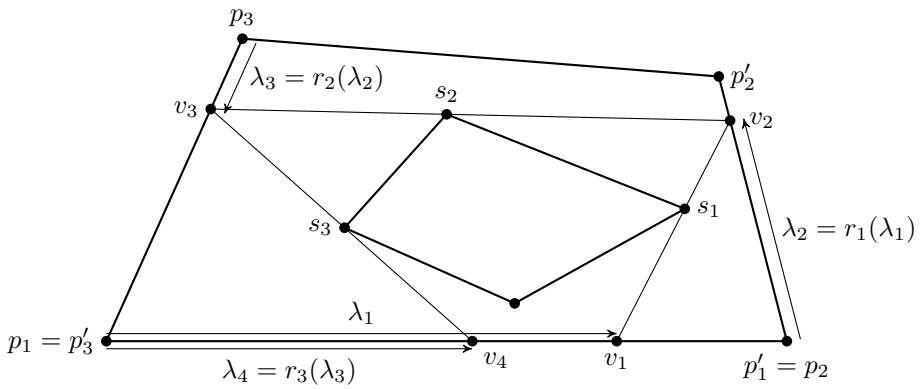

Figure 2 Supporting polygon Q_{v_1}. For every $i \in [3]$, vertex v_i lies on edge $p_i p_i'$ of P, and s_i is the point where the supporting line segment $v_i v_{i+1}$ touches the inner polygon S on its left.

polygon if and only if $(r_k \circ \ldots \circ r_1)(\lambda) \geq \lambda$, where λ is the convex representation of v_1.

It follows from [1] that, for each edge of P, one can compute in polynomial time a partition $\mathcal{I}$ of $[0,1]$ into intervals with rational endpoints such that if λ_1, λ_2 are in the same interval $I \in \mathcal{I}$ then the points with convex representation λ_1 and λ_2 are associated with the same sequence of ray functions $r_1, \ldots, r_k$. Using Lemma 9 we can, for each interval $I \in \mathcal{I}$, compute these ray functions in polynomial time. Define the *slack function* $s(\lambda) = (r_k \circ \ldots \circ r_1)(\lambda) - \lambda$. In fact, this function has the form $s(\lambda) = \frac{a\lambda+b}{c\lambda+d} - \lambda$ for rational numbers a, b, c, d that are also computable in polynomial time (see Corollary 15 in Appendix C). Then it is straightforward to check whether $s(\lambda) \geq 0$ has a solution $\lambda \in I$.

Next we show that if such a solution exists, then there exists a rational solution, which, moreover, can be computed in polynomial time. To this end, let $\lambda^* \in I$ be such that $s(\lambda^*) \geq 0$. If λ^* is on the boundary of I, then $\lambda^* \in \mathbb{Q}$. If λ^* is not on the boundary and is not an isolated solution, then there exists a rational solution in its neighbourhood. Lastly, let λ^* be an isolated solution not on the boundary. Then, λ^* is a root of both s and its derivative s'. For every $\lambda \in I$, we have

$$(c\lambda + d) \cdot s(\lambda) = a\lambda + b - \lambda \cdot (c\lambda + d).$$

Taking the derivative of the above equation with respect to λ, we get

$$c \cdot s(\lambda) + (c\lambda + d) \cdot s'(\lambda) = a - d - 2c\lambda. \tag{1}$$

Since $s(\lambda^*) = s'(\lambda^*) = 0$, from (1) we get $0 = a - d - 2c\lambda^*$. Note that $c \neq 0$ since otherwise $s \equiv 0$. Therefore, $\lambda^* = \frac{a-d}{2c} \in \mathbb{Q}$.

It follows that the vertex v represented by λ^* has rational coordinates computable in polynomial time. By computing $(r_i \circ \ldots \circ r_1)(\lambda^*)$ for $i \in [k]$, we can compute in polynomial time the convex representation of all vertices of the supporting polygon Q_v. Observe, in particular, that all vertices are rational. Hence we have proved Theorem 8.

5 Restricted NMF Requires Irrationality

Here we show that the restricted nonnegative ranks over $\mathbb{R}$ and $\mathbb{Q}$ are, in general, different.

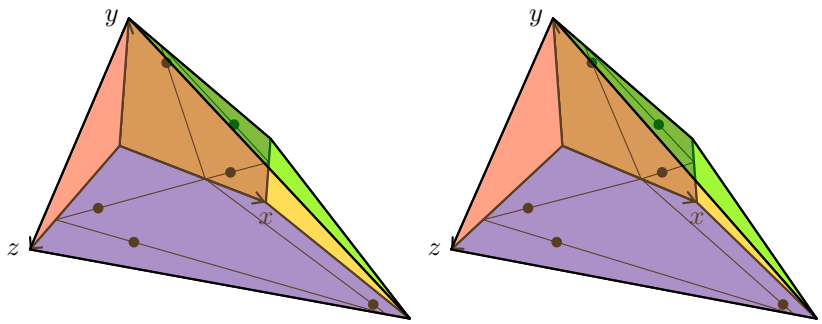

Figure 3 Instance of the nested polytope problem. The two images show orthogonal projections of a 3-dimensional outer polytope P. The black dots indicate 6 inner points (3 on the brown xy-face, and 3 on the blue xz-face) that span the interior polytope S. The two triangles on the xy-face and on the xz-face indicate the (unique) location of 5 points that span the nested polytope Q. The two slightly different projections are designed to create a 3-dimensional impression using stereoscopy. The "parallel-eye" technique should be used, see, e.g., [18]. See Figure 8 in Appendix D for a "cross-eyed" variant.

▶ **Theorem 10.** *Let*

$$M = \begin{pmatrix} 1/8 & 1/2 & 17/22 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1/2 & 3/4 & 7/12 \\ 3/4 & 3/4 & 3/11 & 2 & 1/2 & 1/6 \\ 1/4 & 1/4 & 8/11 & 1/4 & 19/8 & 55/24 \\ 1/2 & 1/8 & 1/11 & 1/8 & 15/16 & 17/16 \\ 11/16 & 5/16 & 7/44 & 1/16 & 7/32 & 43/96 \end{pmatrix} \in \mathbb{Q}_+^{6\times 6}.$$

The restricted nonnegative rank of M over $\mathbb{R}$ is 5. The restricted nonnegative rank of M over $\mathbb{Q}$ is 6.

Proof. Matrix M has an NMF $M = W \cdot H$ with inner dimension 5 with W, H as follows:

$$W = \begin{pmatrix} 0 & \frac{3+\sqrt{2}}{14} & \frac{11+\sqrt{2}}{14} & 0 & 0 \\ 0 & 0 & 0 & \frac{12-2\sqrt{2}}{17} & \frac{5}{7}+\frac{\sqrt{2}}{14} \\ 2-\sqrt{2} & 1 & \frac{3-\sqrt{2}}{7} & \frac{26+7\sqrt{2}}{17} & 0 \\ -1+\sqrt{2} & 0 & \frac{4+\sqrt{2}}{7} & \frac{21-12\sqrt{2}}{17} & \frac{39}{14}+\frac{5\sqrt{2}}{28} \\ \frac{\sqrt{2}}{2} & \frac{2}{7}-\frac{\sqrt{2}}{14} & 0 & \frac{7-4\sqrt{2}}{17} & \frac{33}{28}+\frac{\sqrt{2}}{56} \\ \frac{1}{2}+\frac{\sqrt{2}}{4} & \frac{15}{28}-\frac{\sqrt{2}}{14} & \frac{3-\sqrt{2}}{28} & 0 & \frac{3}{8}-\frac{\sqrt{2}}{16} \end{pmatrix},$$

$$H = \begin{pmatrix} \frac{1+\sqrt{2}}{4} & 0 & \frac{\sqrt{2}}{11} & \frac{1}{4}-\frac{\sqrt{2}}{8} & 0 & \frac{1}{6}+\frac{\sqrt{2}}{12} \\ \frac{3-\sqrt{2}}{4} & \frac{1}{2}+\frac{\sqrt{2}}{8} & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{2}-\frac{\sqrt{2}}{8} & 1-\frac{\sqrt{2}}{11} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{3}{4}+\frac{\sqrt{2}}{8} & \frac{13}{34}-\frac{7\sqrt{2}}{68} & 0 \\ 0 & 0 & 0 & 0 & \frac{21}{34}+\frac{7\sqrt{2}}{68} & \frac{5}{6}-\frac{\sqrt{2}}{12} \end{pmatrix}.$$

Since $\operatorname{rank}(M) = \operatorname{rank}(W) = 4$, the NMF $M = W \cdot H$ is restricted. This RNMF has been obtained by reducing, according to Proposition 2, an NPP instance, which we now describe.

Figure 3 shows the outer 3-dimensional polytope P with 6 faces. The polytope P is the intersection of the following half-spaces: $y \geq 0$ (blue), $z \geq 0$ (brown), $x \geq 0$ (pink), $-x + \frac{5}{2}z + 1 \geq 0$ (yellow), $-\frac{1}{2}x - y + \frac{1}{4}z + 1 \geq 0$ (green), $-\frac{1}{4}x - y - \frac{7}{8}z + 1 \geq 0$

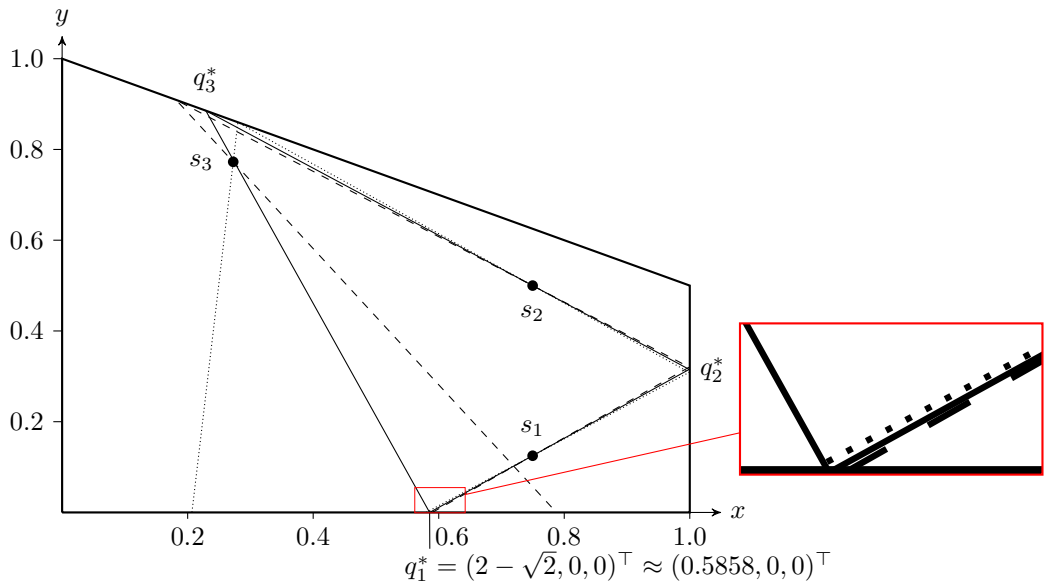

Figure 4 Detailed view of the xy-plane. The outer quadrilateral is one of 6 faces of P, the brown face in Figure 3. The points s_1, s_2, s_3 are among the 6 points that span the inner polytope S. The points q_1^*, q_2^*, q_3^* are among the 5 points that span the nested polytope Q. The area around q_1^* is zoomed in on the right-hand side. The picture illustrates that q_1^* cannot be moved left on the x-axis without increasing the number of vertices of the nested polytope: A dotted ray from a point slightly to the left of q_1^* is drawn through s_1. Its intersection with the line $x = 1$ is slightly below q_2^*. Following the algorithm of [1], the dotted ray is continued in a similar fashion, "wrapping around" s_2 and s_3, and ending on the x-axis at around $x \approx 0.2$, far left of the starting point. On the other hand, the dashed line illustrates that q_1^* could be moved right (considering only this face).

(transparent front). The figure also indicates an interior polytope S spanned by 6 points (black dots): $s_1 = (\frac{3}{4}, \frac{1}{8}, 0)^\top$, $s_2 = (\frac{3}{4}, \frac{1}{2}, 0)^\top$, $s_3 = (\frac{3}{11}, \frac{17}{22}, 0)^\top$, $s_4 = (2, 0, \frac{1}{2})^\top$, $s_5 = (\frac{1}{2}, 0, \frac{3}{4})^\top$, $s_6 = (\frac{1}{6}, 0, \frac{7}{12})^\top$. In the following we discuss possible locations of 5 points q_1, q_2, q_3, q_4, q_5 that span a nested polytope Q. Since s_1, s_2, s_3 all lie on the (brown) face on the xy-plane, but not on a common line, at least 3 of the q_i must lie on the xy-plane. A similar statement holds for s_4, s_5, s_6 and the xz-plane. So at least one q_i, say q_1, must lie on the x-axis.

Suppose another q_i, say q_2, lies on the x-axis. Without loss of generality we can take $q_1 = (0,0,0)^\top$ and $q_2 = (1,0,0)^\top$, as all points in P on the x-axis are enclosed by these q_1, q_2. Figure 4 provides a detailed view of the xy-plane. To enclose s_2, some $q \in \{q_3, q_4, q_5\}$ must also lie on the xy-plane and to the right of the line that connects $q_2 = (1,0,0)^\top$ and s_2. To enclose s_3, some $q' \in \{q_3, q_4, q_5\}$ must also lie on the xy-plane and to the left of the line that connects $q_1 = (0,0,0)^\top$ and s_3. If q and q' were identical then they would lie outside P—a contradiction. Hence 4 points (namely, q_1, q_2, q, q') are on the xy-plane. This leaves only one point, say q'', that is not on the xy-plane. To enclose s_4 (see Figure 9 in Appendix D), point q'' must lie on the xz-plane and must lie to the right of the line that connects $q_2 = (1,0,0)^\top$ and s_4. To enclose s_6, point q'' must lie to the left of the line that connects $q_1 = (0,0,0)^\top$ and s_6. Hence q'' lies outside P—a contradiction.

Hence we have shown that only one point, say q_1, lies on the x-axis, and two points besides q_1, say q_2, q_3, lie on the xy-plane, and two points besides q_1, say q_4, q_5, lie on the xz-plane. Figure 4 indicates a possible location (q_1^*, q_2^*, q_3^*) of q_1, q_2, q_3. The figure illustrates

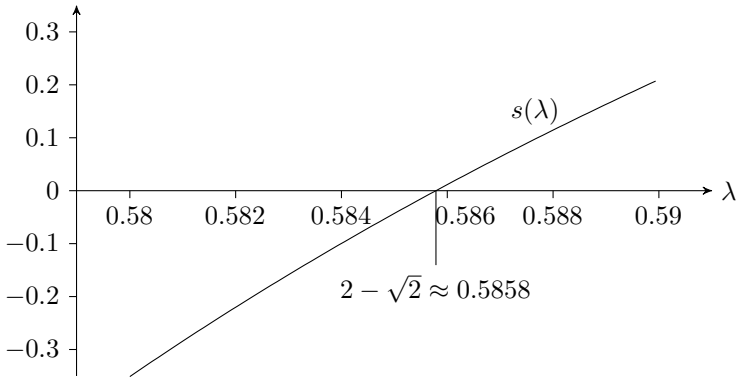

Figure 5 The slack function $s(\lambda) = \frac{52\lambda-30}{15\lambda-8} - \lambda$ corresponding to Figure 4. When $s(\lambda) < 0$, there is no nested triangle with vertex $(\lambda, 0, 0)$.

that the x-coordinate of q_1^* must be at least $2-\sqrt{2}$. Figure 5 illustrates how to prove the same fact more formally, using the concept of a slack function (see Section 4): The slack function $s(\lambda)$ for the interval containing $2-\sqrt{2}$ has a zero at $\lambda = 2-\sqrt{2}$, with a sign change from negative to positive. An inspection of the intervals (of the partition $\mathcal{I}$ from Section 4) to the "left" of $2-\sqrt{2}$ reveals that none of the corresponding slack functions $\tilde{s}$ satisfies $\tilde{s}(\lambda) \geq 0$ for $\lambda < 2-\sqrt{2}$. Similarly, the x-coordinate of q_1^* must be *at most* $2-\sqrt{2}$, as illustrated by Figures 9 and 10 in Appendix D. Hence $q_1^* = (2-\sqrt{2}, 0, 0)^\top$ is necessary. This uniquely (up to permutations) determines q_2^*, q_3^* and similarly the locations q_4^*, q_5^* of q_4, q_5. With the reduction from Proposition 2 this NPP solution determines the RNMF of M mentioned at the beginning of the proof. Since there is no 4-point solution of the NPP instance, we have $\operatorname{rrank}_+(M) = 5$. (Since $\operatorname{rank}(M) = 4$, Lemma 1 implies $\operatorname{rank}_+(M) = 5$.) Since there is no 5-point rational solution of the NPP instance, the restricted nonnegative rank of M over $\mathbb{Q}$ is 6. ◀

6 Conclusion and Future Work

We have shown that an optimal *restricted* nonnegative factorization of a rational matrix may require factors that have irrational entries. An outstanding open problem is whether the same holds for general nonnegative factorizations. An answer to this question will likely shed light on the issue of whether the nonnegative rank can be computed in NP.

Another contribution of the paper has been to develop connections between nonnegative matrix factorization and probabilistic automata, thereby answering an old question concerning the latter. Pursuing this connection, and closely related to the above-mentioned open problem, one can ask whether, given a probabilistic automaton with rational transition probabilities, one can always find a minimal equivalent probabilistic automaton that also has rational transition probabilities.

Acknowledgements. The authors would like to thank Michael Benedikt for stimulating discussions, and anonymous referees for their helpful suggestions. Chistikov was sponsored in part by the ERC Synergy award ImPACT. Kiefer is supported by a University Research Fellowship of the Royal Society. Kiefer, Marušić, Shirmohammadi, and Worrell gratefully acknowledge the support of the EPSRC.

References

1 A. Aggarwal, H. Booth, J. O'Rourke, S. Suri, and C. K. Yap. Finding minimal convex nested polygons. *Information and Computation*, 83(1):98–110, 1989.

2 S. Arora, R. Ge, R. Kannan, and A. Moitra. Computing a nonnegative matrix factorization - provably. In *Proceedings of the 44th Symposium on Theory of Computing (STOC)*, pages 145–162, 2012.

3 F. Bancilhon. A geometric model for stochastic automata. *IEEE Trans. Computers*, 23(12):1290–1299, 1974.

4 M. W. Berry, N. Gillis, and F. Glineur. Document classification using nonnegative matrix factorization and underapproximation. In *International Symposium on Circuits and Systems (ISCAS)*, pages 2782–2785. IEEE, 2009.

5 S. S. Bucak and B. Günsel. Video content representation by incremental non-negative matrix factorization. In *Proceedings of the International Conference on Image Processing (ICIP)*, pages 113–116. IEEE, 2007.

6 J. Canny. Some algebraic and geometric computations in PSPACE. In *Proceedings of the 20th Annual Symposium on Theory of Computing (STOC)*, pages 460–467, 1988.

7 A. Cichocki, R. Zdunek, A. H. Phan, and S.-i. Amari. *Nonnegative Matrix and Tensor Factorizations: Applications to Exploratory Multi-way Data Analysis and Blind Source Separation*. Wiley Publishing, 2009.

8 J. E. Cohen and U. G. Rothblum. Nonnegative ranks, decompositions, and factorizations of nonnegative matrices. *Linear Algebra and its Applications*, 190:149–168, 1993.

9 D. L. Donoho and V. Stodden. When does non-negative matrix factorization give a correct decomposition into parts? In *Advances in Neural Information Processing Systems (NIPS)*, pages 1141–1148, 2003.

10 N. Fijalkow, S. Kiefer, and M. Shirmohammadi. Trace refinement in labelled Markov decision processes. In *Proceedings of the 19th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS)*, volume 9634 of *Lecture Notes in Computer Science*, pages 303–318. Springer, 2016.

11 N. Gillis and F. Glineur. On the geometric interpretation of the nonnegative rank. *Linear Algebra and its Applications*, 437(11):2685–2712, 2012.

12 N. Gillis and S. A. Vavasis. Semidefinite programming based preconditioning for more robust near-separable nonnegative matrix factorization. *SIAM Journal on Optimization*, 25(1):677–698, 2015.

13 A. Kumar, V. Sindhwani, and P. Kambadur. Fast conical hull algorithms for near-separable non-negative matrix factorization. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, page 231–239, 2013.

14 D. D. Lee and H. S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999.

15 A. Paz. *Introduction to probabilistic automata*. Academic Press, New York, 1971.

16 J. Renegar. On the computational complexity and geometry of the first-order theory of the reals. Parts I–III. *Journal of Symbolic Computation*, 13(3):255–352, 1992.

17 Y. Shitov. Nonnegative rank depends on the field. Technical report, arxiv.org, 2015. Available at `http://arxiv.org/abs/1505.01893`.

18 D. Simanek. How to view 3D without glasses. `https://www.lhup.edu/~dsimanek/3d/view3d.htm`. Online, accessed in April 2016.

19 E. Tjioe, M. W. Berry, and R. Homayouni. Discovering gene functional relationships using FAUN (feature annotation using nonnegative matrix factorization). *BMC Bioinformatics*, 11(S-6):S14, 2010.

20 S. A. Vavasis. On the complexity of nonnegative matrix factorization. *SIAM Journal on Optimization*, 20(3):1364–1377, 2009.

21 T. Yokota, R. Zdunek, A. Cichocki, and Y. Yamashita. Smooth nonnegative matrix and tensor factorizations for robust multi-way data analysis. *Signal Processing*, 113:234–249, 2015.

22 S. Zhang, W. Wang, J. Ford, and F. Makedon. Learning from incomplete ratings using nonnegative matrix factorization. In *Proceedings of the 6th SIAM International Conference on Data Mining*, pages 549–553. SIAM, 2006.

A Proofs of Section 2

We show point 2 from the description of Proposition 2. Let $A \in \mathbb{Q}^{n\times(r-1)}$ and $b \in \mathbb{Q}^n$ such that $P = \{\, x \in \mathbb{R}^{r-1} \mid Ax + b \geq 0 \,\}$ is a full-dimensional polytope. Hence $(A \;\; b) \in \mathbb{Q}^{n\times r}$ has full rank r. Let $S = \mathrm{conv}\{s_1, \ldots, s_m\} \subset P$ be a full-dimensional polytope. Define matrix $M \in \mathbb{Q}^{n\times m}$ with $M^j = As_j + b$ for $j \in [m]$. We show the correspondences (a) and (b) from the main text.

(a) Consider an RNMF $M = W \cdot H$ with inner dimension d. We can assume that W has no zero-columns. We have $\mathrm{Col}(W) = \mathrm{Col}(M) = \mathrm{Col}((A \;\; b))$: indeed, the first equality holds because the NMF is restricted, and the second equality holds as S is full-dimensional. So there is a matrix $C \in \mathbb{R}^{r\times d}$ such that $W = (A \;\; b) \cdot C$. For all $i \in [d]$ define $\hat{c}_i \in \mathbb{R}^{r-1}$ so that $C^i = \begin{pmatrix} \hat{c}_i \\ C_{r,i} \end{pmatrix}$. Since W is nonnegative, we have $A\hat{c}_i + C_{r,i}b \geq 0$. Observe that there is no $y \in \mathbb{R}^{r-1} \setminus \{0\}$ with $Ay \geq 0$: indeed, if $Ay \geq 0$ for some nonzero y then for any $x \in P$ and any $t > 0$ we have $A(x + ty) + b \geq Ax + b \geq 0$, implying that P is unbounded, which is false. We use this observation to show that $C_{r,i} > 0$.

- Towards a contradiction, suppose $C_{r,i} = 0$. Then $A\hat{c}_i \geq 0$. Since W has no zero-columns, we have $\hat{c}_i \neq 0$, contradicting the observation above.
- Towards a contradiction, suppose $C_{r,i} < 0$. By dividing the inequality $A\hat{c}_i + C_{r,i}b \geq 0$ by $-C_{r,i}$ we obtain $A(-\hat{c}_i/C_{r,i}) - b \geq 0$. Let $x \in P \setminus \{\hat{c}_i/C_{r,i}\}$. Then $Ax + b \geq 0$, and by adding the previous inequality, we obtain $A(x - \hat{c}_i/C_{r,i}) \geq 0$. Since $x \neq \hat{c}_i/C_{r,i}$, this also contradicts the observation above.

Thus we have shown that $C_{r,i} > 0$ for all $i \in [d]$. Define a diagonal matrix $D \in \mathbb{R}^{d\times d}$ such that $D_{i,i} = C_{r_i} > 0$, and define $H' = D \cdot H$. Then we have:

$$(A \;\; b)\begin{pmatrix} s_1 & s_2 & \cdots & s_m \\ 1 & 1 & \cdots & 1 \end{pmatrix} = M = W \cdot H = (A \;\; b)\, C \cdot H = (A \;\; b)\, CD^{-1} \cdot DH$$
$$= (A \;\; b)\begin{pmatrix} \hat{c}_1/C_{r,1} & \hat{c}_2/C_{r,2} & \cdots & \hat{c}_d/C_{r,d} \\ 1 & 1 & \cdots & 1 \end{pmatrix} H'.$$

Since the columns of $(A \;\; b)$ are linearly independent, it follows:

$$\begin{pmatrix} s_1 & s_2 & \cdots & s_m \\ 1 & 1 & \cdots & 1 \end{pmatrix} = \begin{pmatrix} \hat{c}_1/C_{r,1} & \hat{c}_2/C_{r,2} & \cdots & \hat{c}_d/C_{r,d} \\ 1 & 1 & \cdots & 1 \end{pmatrix} H'.$$

Considering the last row, we see that each column of H' sums up to 1. Since H' is nonnegative, H' is column-stochastic. For $i \in [d]$ define $q_i = \hat{c}_i/C_{r,i}$. Then for all $j \in [m]$ we have $s_j \in \mathrm{conv}\{q_1, \cdots, q_d\}$. Hence, defining the polytope $Q = \mathrm{conv}\{q_1, \cdots, q_d\}$ we have $S \subseteq Q$. For all $i \in [d]$ we have $A\hat{c}_i + C_{r,i}b \geq 0$, hence $Aq_i + b \geq 0$. It follows $Q \subseteq P$. Thus Q is nested between P and S.

(b) Consider d points $q_1, \ldots, q_d \in \mathbb{R}^{r-1}$ with $S \subseteq \mathrm{conv}\{q_1, \ldots, q_d\} \subseteq P$. Define a matrix $W \in \mathbb{Q}^{n\times d}$ by $W^i = Aq_i + b$ for $i \in [d]$. Matrix W is nonnegative, as $q_i \in P$. Define a column-stochastic matrix $H \in \mathbb{Q}^{d\times m}$ so that $s_j = \sum_{i=1}^{d} H_{i,j}q_i$ for $j \in [m]$. Such $H_{i,j}$

exist, as $s_j \in \operatorname{conv}\{q_1, \ldots, q_d\}$. We have for all $j \in [m]$:

$$\begin{aligned}
M^j &= As_j + b && \text{definition of } M \\
&= A\left(\sum_{i=1}^{d} H_{i,j} q_i\right) + b && \text{definition of } H \\
&= \left(\sum_{i=1}^{d} Aq_i \cdot H_{i,j}\right) + b \\
&= \sum_{i=1}^{d} (\underbrace{Aq_i + b}_{W^i}) \cdot H_{i,j} && \text{as } H \text{ is column-stochastic} \\
&= (W \cdot H)^j .
\end{aligned}$$

Hence $M = W \cdot H$ is an NMF. Since S is full-dimensional, we have $\operatorname{Col}(M) = \operatorname{Col}((A \;\; b))$. As $\operatorname{Col}(W) \subseteq \operatorname{Col}((A \;\; b))$ it follows that $\operatorname{Col}(W) \subseteq \operatorname{Col}(M)$, hence NMF $M = W \cdot H$ is restricted.

B Proofs of Section 3

B.1 Discussion of Erroneous Claims in [3]

As mentioned in the main text, Bancilhon [3, Theorem 13] claims a statement that implies a positive answer to Paz's Question 4. We have shown that the correct answer to Paz's question is negative. In the following we track down where the paper [3] goes wrong. The proof of [3, Theorem 13] offered therein relies on another (wrong) claim about *cones*.

Let $\mathcal{V}$ be a vector space. Let $v_1, \ldots, v_n \in \mathcal{V}$. The *polyhedral cone generated by* $v_1, \ldots, v_n$ is the set

$$\left\{ \sum_{i=1}^{n} \lambda_i v_i \;\middle|\; \lambda_1, \ldots, \lambda_n \geq 0 \right\} \subseteq \mathcal{V} .$$

▶ **Claim 11** ([3, Theorem 2], slightly paraphrased)**.** *Let $\mathcal{V}$ be a vector space. Let $\mathcal{V}' \subseteq \mathcal{V}$ be a subspace of $\mathcal{V}$. Let $\mathcal{C} \subseteq \mathcal{V}$ be a polyhedral cone generated by n vectors that also span $\mathcal{V}$. Then $\mathcal{C} \cap \mathcal{V}'$ is a polyhedral cone generated by at most n vectors.*

Claim 11 is false. For a counterexample, consider the polyhedral cone $\mathcal{C}$ generated by the following 6 vectors:

$$\begin{pmatrix} 1\\0\\0\\1\\0\\0 \end{pmatrix}, \quad \begin{pmatrix} 0\\0\\0\\1\\0\\0 \end{pmatrix}, \quad \begin{pmatrix} 0\\1\\0\\0\\1\\0 \end{pmatrix}, \quad \begin{pmatrix} 0\\0\\0\\0\\1\\0 \end{pmatrix}, \quad \begin{pmatrix} 0\\0\\1\\0\\0\\1 \end{pmatrix}, \quad \begin{pmatrix} 0\\0\\0\\0\\0\\1 \end{pmatrix}$$

It can be described equivalently by the conjunction of inequalities $0 \leq x_1 \leq x_4$ and $0 \leq x_2 \leq x_5$ and $0 \leq x_3 \leq x_6$. Let $\mathcal{V}' \subset \mathbb{R}^6$ be the vector space defined by the equalities $x_4 = x_5 = x_6$. Then the cone $\mathcal{C}' = \mathcal{C} \cap \mathcal{V}'$ can be described by the inequalities $0 \leq x_1, x_2, x_3 \leq x_4 = x_5 = x_6$.

The cone $\mathcal{C}'$ is generated by the following 8 vectors:

$$\begin{pmatrix}0\\0\\0\\1\\1\\1\end{pmatrix}, \begin{pmatrix}0\\0\\1\\1\\1\\1\end{pmatrix}, \begin{pmatrix}0\\1\\0\\1\\1\\1\end{pmatrix}, \begin{pmatrix}0\\1\\1\\1\\1\\1\end{pmatrix}, \begin{pmatrix}1\\0\\0\\1\\1\\1\end{pmatrix}, \begin{pmatrix}1\\0\\1\\1\\1\\1\end{pmatrix}, \begin{pmatrix}1\\1\\0\\1\\1\\1\end{pmatrix}, \begin{pmatrix}1\\1\\1\\1\\1\\1\end{pmatrix}.$$

All those vectors are extremal in cone $\mathcal{C}'$, so it cannot be generated by fewer than 8 vectors. Hence, Claim 11 is false.

Let us further examine how Claim 11 was justified in [3]. The proof offered therein starts with the following claim, which is stated there without further justification:

▶ **Claim 12** (proof of [3, Theorem 2], slightly paraphrased). *A cone $\mathcal{C}$ is generated by n vectors if and only if it is limited by n hyperplanes.*

Claim 12 is also false. For a counterexample, consider the cone $\mathcal{C} \subseteq \mathbb{R}^4_+$ limited by the following 6 hyperplanes:

$$0 \le x_1 \le y, \quad 0 \le x_2 \le y, \quad 0 \le x_3 \le y.$$

(All vectors in $\mathcal{C}$ with $y = y^*$ form a cube of length y^*.) The following set contains 8 vectors, all of which are extremal in $\mathcal{C}$:

$$\{0,1\} \times \{0,1\} \times \{0,1\} \times \{1\} \quad \subset \quad \mathcal{C}.$$

Hence $\mathcal{C}$ is not generated by 6 vectors, contradicting Claim 12.

B.2 Details of the Proof of Proposition 6

In this subsection, we complete some details from the proof of Proposition 6.

Proposition 6. *Given a nonnegative matrix $M \in \mathbb{Q}_+^{n\times m}$ of rank r, one can compute in polynomial time an LMC $\mathcal{M} = (m+2, \Sigma, \mu)$ of rank $r+2$ such that for all $d \in \mathbb{N}$:*

(a) any d-dimensional NMF $M = W \cdot H$ determines an LMC $\mathcal{M}' = (d+2, \Sigma, \mu')$ with $\mathcal{M}' \ge \mathcal{M}$ and $\mathrm{rank}(\mathcal{M}') = \mathrm{rank}(W) + 2$, and

(b) any LMC $\mathcal{M}' = (d+2, \Sigma, \mu')$ with $\mathcal{M}' \ge \mathcal{M}$ determines a d-dimensional NMF $M = W \cdot H$ with $\mathrm{rank}(\mathcal{M}') = \mathrm{rank}(W) + 2$.

In particular, for all $d \in \mathbb{N}$ the inequality $\mathrm{rrank}_+(M) \le d$ holds if and only if $\mathcal{M}$ is covered by some $(d+2)$-state LMC $\mathcal{M}'$ such that $\mathcal{M}'$ and $\mathcal{M}$ have the same rank.

Here we will write $\mathbf{0}$ for the column vector with all zeros, whose dimension will be clear from the context. For every $p \in \mathbb{N}$, we denote by $\checkmark^p$ the p-fold concatenation of $\checkmark$ by itself.

Let us first take a detailed look at $\mathrm{Back}\,\mathcal{M}$. For every $i \in [m]$ and $j \in [n]$, we have

$$(\mathrm{Back}\,\mathcal{M})_{i,b_j} = \mu(b_j)_i \cdot \mathbf{1} = \mu(b_j)_{i,m+1} = M_{j,i}.$$

From here it is easy to see that for every $w \in \Sigma^*$:

$$(\operatorname{Back}\mathcal{M})^w = \begin{cases} \mathbf{1} & w = \varepsilon \\ \frac{1}{m} \cdot e_0 & w = a_i \text{ with } i \in [m] \\ e_{m+1} & w = \checkmark^p \text{ with } p \in \mathbb{N} \\ (0, M_j, 0)^\top & w = b_j\checkmark^p \text{ with } j \in [n], p \in \mathbb{N}_0 \\ \left(\frac{1}{m} M_{j,i}\right) \cdot e_0 & w = a_i b_j \checkmark^p \text{ with } i \in [m], j \in [n], p \in \mathbb{N}_0 \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

That is,

$$\operatorname{Back}\mathcal{M} = \begin{array}{c} \begin{array}{ccccccccccccc} \scriptstyle\varepsilon & \scriptstyle b_1 & \scriptstyle\cdots & \scriptstyle b_n & \scriptstyle\checkmark & \scriptstyle a_i & \scriptstyle a_i b_j & \scriptstyle b_1\checkmark & \scriptstyle\cdots & \scriptstyle b_n\checkmark & \scriptstyle\checkmark^2 & \scriptstyle\cdots \end{array} \\ \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 & \frac{1}{m} & \frac{1}{m}M_{j,i} & 0 & \cdots & 0 & 0 & \cdots \\ 1 & M_{1,1} & \cdots & M_{n,1} & 0 & 0 & 0 & M_{1,1} & \cdots & M_{n,1} & 0 & \cdots \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \cdots \\ 1 & M_{1,m} & \cdots & M_{n,m} & 0 & 0 & 0 & M_{1,m} & \cdots & M_{n,m} & 0 & \cdots \\ 1 & 0 & \cdots & 0 & 1 & 0 & 0 & 0 & \cdots & 0 & 1 & \cdots \end{pmatrix} \end{array}$$

Clearly, the following submatrix of $\operatorname{Back}\mathcal{M}$ has the same column space as $\operatorname{Back}\mathcal{M}$:

$$\begin{array}{c} \\ \scriptstyle 0 \\ \scriptstyle 1 \\ \scriptstyle\vdots \\ \scriptstyle m \\ \scriptstyle m{+}1 \end{array} \begin{array}{c} \begin{array}{ccccc} \scriptstyle\varepsilon & \scriptstyle b_1 & \scriptstyle\cdots & \scriptstyle b_n & \scriptstyle\checkmark \end{array} \\ \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 \\ 1 & M_{1,1} & \cdots & M_{n,1} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & M_{1,m} & \cdots & M_{n,m} & 0 \\ 1 & 0 & \cdots & 0 & 1 \end{pmatrix} \end{array}$$

This implies that $\operatorname{rank}(\mathcal{M}) = \operatorname{rank}(M) + 2 = r + 2$.

For direction (a), consider an NMF $M = W \cdot H$ and the LMC $\mathcal{M}'$ defined in the main text. For its backward matrix $\operatorname{Back}\mathcal{M}' \in \mathbb{R}_+^{\{0,1,\ldots,d,d+1\}\times\Sigma^*}$, it is easy to see from the definition that for every $w \in \Sigma^*$:

$$(\operatorname{Back}\mathcal{M}')^w = \begin{cases} \mathbf{1} & w = \varepsilon \\ \frac{1}{m} \cdot e_0 & w = a_i \text{ with } i \in [m] \\ e_{d+1} & w = \checkmark^p \text{ with } p \in \mathbb{N} \\ (0, W_j, 0)^\top & w = b_j\checkmark^p \text{ with } j \in [n], p \in \mathbb{N}_0 \\ \left(\frac{1}{m} M_{j,i}\right) \cdot e_0 & w = a_i b_j \checkmark^p \text{ with } i \in [m], j \in [n], p \in \mathbb{N}_0 \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

Indeed, for every $i \in [m]$:

$$(\operatorname{Back}\mathcal{M}')^{a_i} = \left(\sum_{l \in [d]} \frac{1}{m} H_{l,i}\right) \cdot e_0 = \frac{1}{m} \cdot e_0,$$

since H is a column-stochastic matrix. From here it is clear that the columns of the following

submatrix of $\operatorname{Back}\mathcal{M}'$ span $\operatorname{Col}(\operatorname{Back}\mathcal{M}')$:

$$\begin{array}{c} \\ 0 \\ 1 \\ \vdots \\ d \\ d{+}1 \end{array} \begin{array}{c} \begin{array}{cccccc} \varepsilon & b_1 & \cdots & b_n & \checkmark \end{array} \\ \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 \\ 1 & W_{1,1} & \cdots & W_{n,1} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & W_{1,d} & \cdots & W_{n,d} & 0 \\ 1 & 0 & \cdots & 0 & 1 \end{pmatrix} \end{array}$$

Hence, $\operatorname{rank}(\mathcal{M}') = \operatorname{rank}(W) + 2$. In particular, if NMF $M = W \cdot H$ is restricted, then $\operatorname{rank}(\mathcal{M}) = \operatorname{rank}(M) + 2 = \operatorname{rank}(W) + 2 = \operatorname{rank}(\mathcal{M}')$.

Now we show in more detail that $\mathcal{M}' \geq \mathcal{M}$. We define a row-stochastic matrix:

$$A = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 \\ 0 & H_{1,1} & \cdots & H_{d,1} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & H_{1,m} & \cdots & H_{d,m} & 0 \\ 0 & 0 & \cdots & 0 & 1 \end{pmatrix} \in \mathbb{R}_+^{\{0,1,\ldots,m,m+1\}\times\{0,1,\ldots,d,d+1\}}.$$

For every $i \in [m]$ and $j \in [n]$:

$$(A \cdot \operatorname{Back}\mathcal{M}')_{i,b_j} = A_i \cdot (\operatorname{Back}\mathcal{M}')^{b_j} = (H^i)^\top \cdot (W_j)^\top = (W_j \cdot H^i)^\top = M_{j,i} = (\operatorname{Back}\mathcal{M})_{i,b_j}.$$

From here it follows easily that $A \cdot \operatorname{Back}\mathcal{M}' = \operatorname{Back}\mathcal{M}$. Hence, $\mathcal{M}' \geq \mathcal{M}$.

Lastly we prove the other direction, (b). Let LMC $\mathcal{M}' = (d+2, \Sigma, \mu')$ be such that $\mathcal{M}' \geq \mathcal{M}$. This implies that there exists a row-stochastic matrix A such that $A \cdot \operatorname{Back}\mathcal{M}' = \operatorname{Back}\mathcal{M}$. Given nonempty subsets I and J of the row and column indices of a matrix C, respectively, we write $C_{I,J}$ for the submatrix $(C_{i,j})_{i\in I, j\in J}$. It is clear from the description of $\operatorname{Back}\mathcal{M}$ that M has the following $(d+2)$-dimensional nonnegative factorization:

$$M^\top = (\operatorname{Back}\mathcal{M})_{[m],\{b_1,\ldots,b_n\}} = A_{[m],\{0,1,\ldots,d,d+1\}} \cdot (\operatorname{Back}\mathcal{M}')_{\{0,1,\ldots,d,d+1\},\{b_1,\ldots,b_n\}}. \quad (2)$$

Assuming M is a non-zero matrix, there exist $i \in [m]$ and $j \in [n]$ such that $M_{j,i} > 0$. Then,

$$\frac{1}{m} M_{j,i} = (\operatorname{Back}\mathcal{M})_{0,a_ib_j} = A_0 \cdot (\operatorname{Back}\mathcal{M}')^{a_ib_j} = \sum_{i_1=0}^{d+1} A_{0,i_1} \cdot (\operatorname{Back}\mathcal{M}')_{i_1,a_ib_j} > 0.$$

Since A and $(\operatorname{Back}\mathcal{M}')$ are nonnegative matrices, there exists $i_1 \in \{0,1,\ldots,d,d+1\}$ such that $A_{0,i_1} \cdot (\operatorname{Back}\mathcal{M}')_{i_1,a_ib_j} > 0$. Without loss of generality, we may assume that $i_1 = 0$. That is,

$$A_{0,0} \cdot (\operatorname{Back}\mathcal{M}')_{0,a_ib_j} > 0. \quad (3)$$

Moreover, $(\operatorname{Back}\mathcal{M})_{m+1,\checkmark} = A_{m+1} \cdot (\operatorname{Back}\mathcal{M}')^{\checkmark} = 1$. The nonnegativity implies that there exists $i_2 \in \{0,1,\ldots,d,d+1\}$ such that $A_{m+1,i_2} \cdot (\operatorname{Back}\mathcal{M}')_{i_2,\checkmark} > 0$. Note that $i_2 \neq 0$ since otherwise we would have that $(\operatorname{Back}\mathcal{M})_{m+1,a_ib_j} \geq A_{m+1,0} \cdot (\operatorname{Back}\mathcal{M}')_{0,a_ib_j} > 0$ which is a contradiction. We may therefore, without loss of generality, assume that $i_2 = d+1$, i.e.,

$$A_{m+1,d+1} \cdot (\operatorname{Back}\mathcal{M}')_{d+1,\checkmark} > 0. \quad (4)$$

▶ **Lemma 13.** *It holds that* $M^\top = A_{[m],[d]} \cdot (\operatorname{Back}\mathcal{M}')_{[d],\{b_1,\ldots,b_n\}}$, *where matrix* $A_{[m],[d]}$ *is row-stochastic.*

Proof. By (2), it suffices to show that

$$A_{[m],\{0,1,\ldots,d,d+1\}} \cdot (\operatorname{Back} \mathcal{M}')_{\{0,1,\ldots,d,d+1\},\{b_1,\ldots,b_n\}} = A_{[m],[d]} \cdot (\operatorname{Back} \mathcal{M}')_{[d],\{b_1,\ldots,b_n\}}.$$

That is, for any $l_1 \in [m]$ and $l_2 \in [n]$ we need to show that

$$A_{l_1,\{0,1,\ldots,d,d+1\}} \cdot (\operatorname{Back} \mathcal{M}')_{\{0,1,\ldots,d,d+1\},b_{l_2}} = A_{l_1,[d]} \cdot (\operatorname{Back} \mathcal{M}')_{[d],b_{l_2}}.$$

To do this, it suffices to show that $A_{l_1,0} \cdot (\operatorname{Back} \mathcal{M}')_{0,b_{l_2}} + A_{l_1,d+1} \cdot (\operatorname{Back} \mathcal{M}')_{d+1,b_{l_2}} = 0$. In the following, we prove that $A_{l_1,0} = 0$ and $A_{l_1,d+1} = 0$. By an analogous argument, it also holds that $(\operatorname{Back} \mathcal{M}')_{0,b_{l_2}} = 0$ and $(\operatorname{Back} \mathcal{M}')_{d+1,b_{l_2}} = 0$.

If $A_{l_1,0} > 0$, then from (3) it would follow that

$$(\operatorname{Back} \mathcal{M})_{l_1,a_ib_j} = A_{l_1} \cdot (\operatorname{Back} \mathcal{M}')^{a_ib_j} \geq A_{l_1,0} \cdot (\operatorname{Back} \mathcal{M}')_{0,a_ib_j} > 0,$$

which is a contradiction since $(\operatorname{Back} \mathcal{M})_{l_1,a_ib_j} = 0$ for all $l_1 \in [m]$. Similarly, if $A_{l_1,d+1} > 0$ then by (4) we would have that

$$(\operatorname{Back} \mathcal{M})_{l_1,\checkmark} = A_{l_1} \cdot (\operatorname{Back} \mathcal{M}')^{\checkmark} \geq A_{l_1,d+1} \cdot (\operatorname{Back} \mathcal{M}')_{d+1,\checkmark} > 0,$$

which is a contradiction since $(\operatorname{Back} \mathcal{M})_{l_1,\checkmark} = 0$ for all $l_1 \in [m]$. Hence, $A_{l_1,0} = A_{l_1,d+1} = 0$.

We have thus shown that $A_{[m],\{0\}} = A_{[m],\{d+1\}} = (0,\ldots,0)^\top$. Since A is row stochastic, this implies that $A_{[m],[d]}$ is row-stochastic. ◀

From Lemma 13 we get a d-dimensional nonnegative factorizaton $M = W \cdot H$ where $W = ((\operatorname{Back} \mathcal{M}')_{[d],\{b_1,\ldots,b_n\}})^\top$ and $H = (A_{[m],[d]})^\top$.

Suppose $\mathcal{M}'$ and $\mathcal{M}$ have the same rank. Then $\operatorname{Row}(\operatorname{Back} \mathcal{M}) = \operatorname{Row}(\operatorname{Back} \mathcal{M}')$, as we already knew that $\operatorname{Row}(\operatorname{Back} \mathcal{M}) \subseteq \operatorname{Row}(\operatorname{Back} \mathcal{M}')$ since $A \cdot \operatorname{Back} \mathcal{M}' = \operatorname{Back} \mathcal{M}$. Take any $v \in \operatorname{Col}(W)$. Since $W^\top = (\operatorname{Back} \mathcal{M}')_{[d],\{b_1,\ldots,b_n\}}$, there exists $u \in \operatorname{Row}(\operatorname{Back} \mathcal{M}')$ such that $v^\top = u_{\{b_1,\ldots,b_n\}}$. As $\operatorname{Row}(\operatorname{Back} \mathcal{M}') = \operatorname{Row}(\operatorname{Back} \mathcal{M})$, from $\operatorname{Back} \mathcal{M}$ it is clear that $v \in \operatorname{Col}(M)$. Hence, $\operatorname{Col}(W) \subseteq \operatorname{Col}(M)$ and therefore $\operatorname{rank}(W) \leq \operatorname{rank}(M)$. Since $M = W \cdot H$, this implies that $\operatorname{rank}(W) = \operatorname{rank}(M)$ as required. This completes the proof of Proposition 6.

C Proofs of Section 4

First we prove Lemma 9 from the main text.

Lemma 9. *Consider bounded polygons $S \subseteq P \subseteq \mathbb{R}^2$ whose vertices are rational and of bit-length L. Then the ray function associated with a supporting line segment uv has the form $r(\lambda) = \frac{a\lambda+b}{c\lambda+d}$, where coefficients a, b, c, d are rational numbers with bit-length $O(L)$ that can be computed in polynomial time.*

Proof. Let vertices u and v lie on edges p_1p_1' and p_2p_2' of P, respectively. Let s_1 be the point where the supporting line segment uv touches the inner polygon S on its left. Let λ_1 and λ_2 be the convex representations of u and v, respectively. That is, $u = (1-\lambda_1)p_1 + \lambda_1 p_1'$ and $v = (1-\lambda_2)p_2 + \lambda_2 p_2'$. By definition of the ray function r we have $\lambda_2 = r(\lambda_1)$.

Let us first consider the case when lines p_1p_1' and p_2p_2' are parallel; see Figure 6 for illustration. Let t denote the intersection of lines p_1p_1' and s_1p_2, and let t' denote the intersection of lines p_2p_2' and s_1p_1'. Since $s_1 \in \mathbb{Q}^2$, we have $t = (1-b) \cdot p_1 + b \cdot p_1'$ and $t' = (1-d) \cdot p_2 + d \cdot p_2'$ for some $b, d \in \mathbb{Q}$. Numbers b, d can be computed from the input vertices p_1, p_1', p_2, p_2' using a constant number of arithmetic operations and therefore have

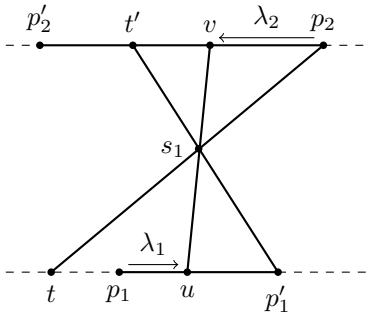

Figure 6 Lines p_1p_1' and p_2p_2' are parallel.

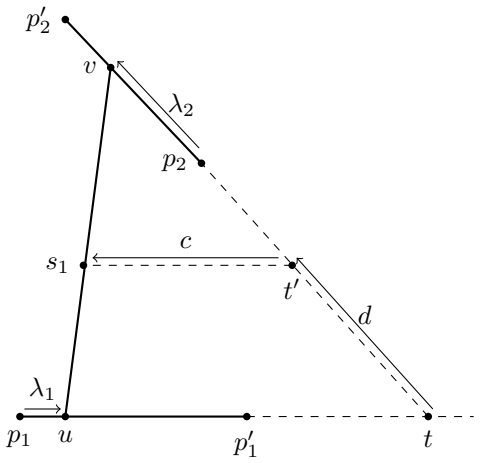

Figure 7 Lines p_1p_1' and p_2p_2' intersect at point t.

bit-length $O(L)$. Since $\triangle tus_1 \sim \triangle p_2vs_1$ and $\triangle tp_1's_1 \sim \triangle p_2t's_1$, we have $\frac{ut}{vp_2} = \frac{ts_1}{s_1p_2}$ and $\frac{ts_1}{s_1p_2} = \frac{p_1't}{t'p_2}$. Hence, $\frac{ut}{vp_2} = \frac{p_1't}{t'p_2}$ and therefore $\frac{\lambda_1 - b}{\lambda_2} = \frac{1-b}{d}$. Hence,

$$\lambda_2 = \frac{d}{1-b} \cdot \lambda_1 - \frac{bd}{1-b}$$

where the coefficients $\frac{d}{1-b}, \frac{bd}{1-b} \in \mathbb{Q}$ have bit-length $O(L)$.

Let us now consider the second case: when lines p_1p_1' and p_2p_2' are not parallel. Let t denote their intersection; see Figure 7 for illustration. Note that $t \in \mathbb{Q}^2$. Without loss of generality, let p_1' be a convex combination of $\{p_1, t\}$ and p_2 be a convex combination of $\{t, p_2'\}$. Then $p_1' = (1-a) \cdot p_1 + a \cdot t$ and $p_2 = (1-b) \cdot t + b \cdot p_2'$ for some $a, b \in [0, 1]$. Numbers a, b are rational as they are each the unique solution of a linear system with rational coefficients. Moreover a, b can be computed from the input vertices p_1, p_1', p_2, p_2' using a constant number of arithmetic operations and therefore have bit-length $O(L)$. Using the above representation of p_1' as a convex combination of $\{p_1, t\}$, we get

$$u = (1-\lambda_1) \cdot p_1 + \lambda_1 \cdot ((1-a) \cdot p_1 + a \cdot t) = (1 - a\lambda_1) \cdot p_1 + a\lambda_1 \cdot t$$

and therefore $u - t = (1 - a\lambda_1) \cdot (p_1 - t)$. Similarly, we have

$$v = (1 - b - \lambda_2 + \lambda_2 b) \cdot t + (b + \lambda_2 - \lambda_2 b) \cdot p_2'$$

and therefore $v - t = (b + \lambda_2 - \lambda_2 b) \cdot (p_2' - t)$. Let the line through s_1 parallel with $p_1 t$ intersect $p_2' t$ at point t'. Note that $s_1 - t' = c \cdot (p_1 - t)$ and $t' - t = d \cdot (p_2' - t)$ for some $c, d \in [0, 1] \cap \mathbb{Q}$. Numbers c, d can be computed from vertices p_1, p_1', p_2, p_2' using a constant number of arithmetic operations and thus have bit-length $O(L)$. Since $\triangle vt's_1 \sim \triangle vtu$, it holds that $\frac{vt'}{vt} = \frac{s_1 t'}{ut}$. We therefore have that

$$\frac{b + \lambda_2 - \lambda_2 b - d}{b + \lambda_2 - \lambda_2 b} = \frac{c}{1 - a\lambda_1}.$$

From the above equation we get that

$$\lambda_2 = \frac{a(b-d) \cdot \lambda_1 + b(c-1) + d}{a(b-1) \cdot \lambda_1 + (b-1)(c-1)}$$

where $a(b-d), b(c-1) + d, a(b-1), (b-1)(c-1) \in \mathbb{Q}$ have bit-length $O(L)$. ◀

A function $f : I \to \mathbb{R}$, where $I \subseteq \mathbb{R}$, is called a *rational linear fractional transformation* if there exist $a, b, c, d \in \mathbb{Q}$ such that $f(\lambda) = \frac{a\lambda + b}{c\lambda + d}$ for every $\lambda \in I$.

▶ **Lemma 14.** *Given two rational linear fractional transformations* $f_1(\lambda) = \frac{a_1\lambda + b_1}{c_1\lambda + d_1}$ *and* $f_2(\lambda) = \frac{a_2\lambda + b_2}{c_2\lambda + d_2}$*, their composition* $f_2 \circ f_1$ *is also a rational linear fractional transformation. Specifically,* $(f_2 \circ f_1)(\lambda) = \frac{a\lambda + b}{c\lambda + d}$ *where* $a, b, c, d \in \mathbb{Q}$ *are such that*

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} a_2 & b_2 \\ c_2 & d_2 \end{bmatrix} \cdot \begin{bmatrix} a_1 & b_1 \\ c_1 & d_1 \end{bmatrix}.$$

Proof. For every λ we have

$$\begin{aligned}(f_2 \circ f_1)(\lambda) &= \frac{a_2 \frac{a_1\lambda + b_1}{c_1\lambda + d_1} + b_2}{c_2 \frac{a_1\lambda + b_1}{c_1\lambda + d_1} + d_2} \\ &= \frac{a_2(a_1\lambda + b_1) + b_2(c_1\lambda + d_1)}{c_2(a_1\lambda + b_1) + d_2(c_1\lambda + d_1)} = \frac{(a_2a_1 + b_2c_1)\lambda + (a_2b_1 + b_2d_1)}{(c_2a_1 + d_2c_1)\lambda + (c_2b_1 + d_2d_1)}.\end{aligned}$$

◀

▶ **Corollary 15.** *Given* $r_i(\lambda) = \frac{a_i\lambda + b_i}{c_i\lambda + d_i}$ *in terms of* $a_i, b_i, c_i, d_i \in \mathbb{Q}$*, for* $i \in [k]$*, one can compute in polynomial time* $a, b, c, d \in \mathbb{Q}$ *such that* $r(\lambda) = (r_k \circ \ldots \circ r_2 \circ r_1)(\lambda) = \frac{a\lambda + b}{c\lambda + d}$.

Proof. By Lemma 14, the coefficients of the composition $r_k \circ \ldots \circ r_2 \circ r_1$ are computable by iterative matrix multiplication:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} a_k & b_k \\ c_k & d_k \end{bmatrix} \cdot \ldots \cdot \begin{bmatrix} a_2 & b_2 \\ c_2 & d_2 \end{bmatrix} \cdot \begin{bmatrix} a_1 & b_1 \\ c_1 & d_1 \end{bmatrix}.$$

The result follows. ◀

D Additional Figures for Section 5

Figure 8 shows a variant of Figure 3, suitable for "cross-eyed" stereoscopic viewing.

Figure 9 shows the xz-plane, in the same way as Figure 4 shows the xy-plane.

Figure 10 plots the corresponding slack function. It shows that a sign change from positive to negative occurs at the root $\lambda = 2 - \sqrt{2}$.

Figure 11 provides a combined view of the xy-plane and the xz-plane, with the coordinates of some vertices of P.

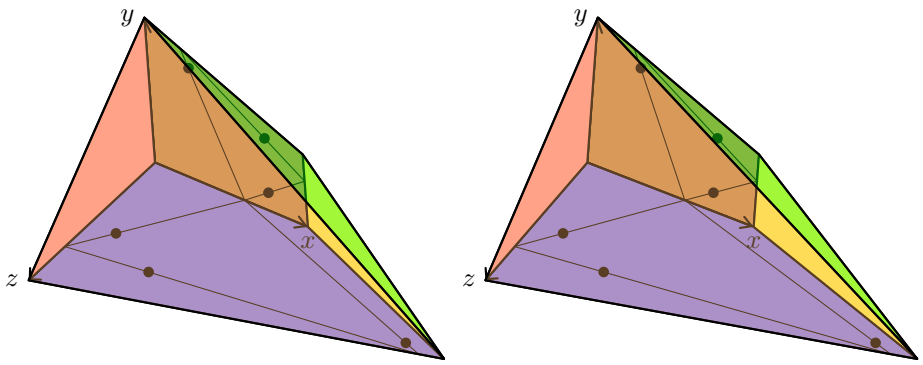

Figure 8 Two projections of P suitable for "cross-eyed" stereoscopic viewing.

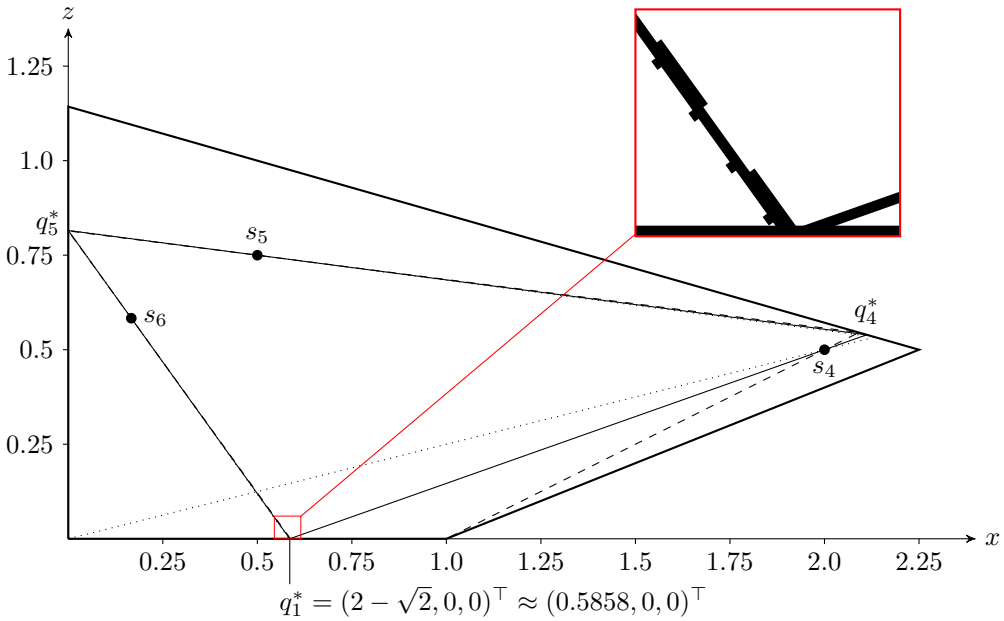

Figure 9 Detailed view on the xz-plane. The image shows that q_1^* cannot be moved right on the x-axis without increasing the number of vertices of the nested polytope.

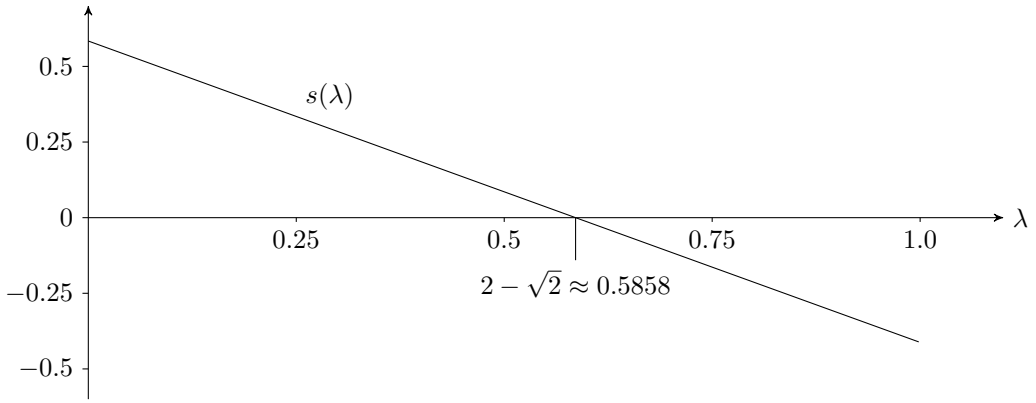

Figure 10 The slack function $s(\lambda) = \frac{23\lambda-80}{40\lambda-137} - \lambda$ corresponding to Figure 9. When $s(\lambda) < 0$, there is no nested triangle with vertex $(\lambda, 0, 0)$.

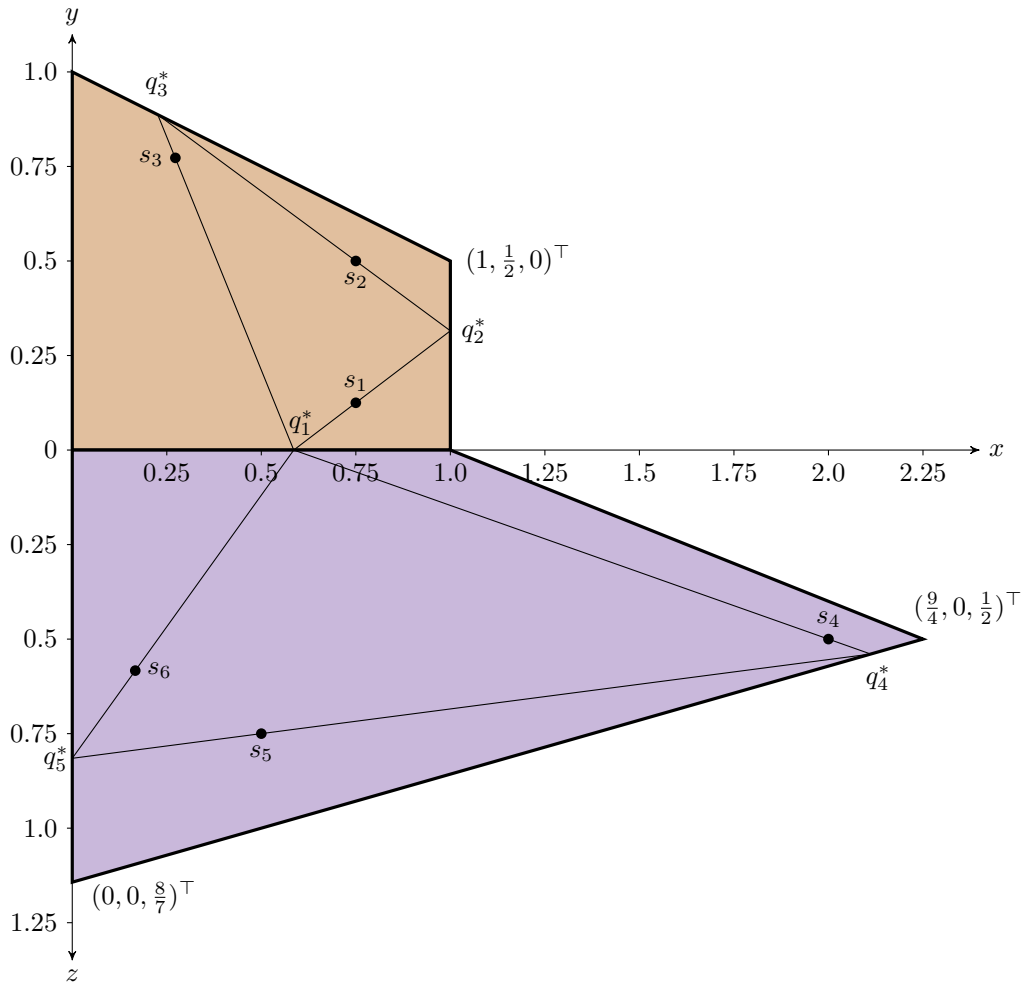

Figure 11 Combined view of the xy-plane and the xz-plane. One may imagine that the figure is folded along the x-axis.